Edge-of-the-Envelope Speculation: Fixed Computational Geographies in Neural Networks via Parameter Atlases and Direct Feedback Alignment

Introduction

The ideas presented here are intentionally speculative and should be regarded as *edge-of-the-envelope hypotheses* rather than established results.

Modern deep learning typically assumes that both:

1. computational routing, and
2. credit assignment,

must be learned and adapt continuously during training.

This note explores the opposite possibility:

> Neural networks may require only a stable computational geography and local adaptive function learning.

Specifically, we consider combining:

- a fixed, locality-preserving **Parameter Atlas** for forward routing; and
- fixed random feedback pathways as in Direct Feedback Alignment (DFA) for backward credit assignment.

Only the local forward parameters are learned.

Conventional Deep Learning View

Standard deep networks simultaneously learn:

Input
↓
Learned routing

↓
Learned representations
↓
Output

Error
↓
Exact backpropagation
↓
Layerwise credit assignment

Routing and learning are tightly coupled.

Proposed Architecture

Forward Pass: Parameter Atlas

The input is transformed into a global locality-preserving routing code (e.g., via locality-sensitive hashing).

Input
↓
Parameter Atlas
(fixed)
↓
Sparse route selection
↓
Local computation

The atlas determines, at the outset, which parameters participate throughout the network.

Routing is:

- deterministic;
- locality preserving;
- fixed after initialization.

Nearby inputs tend to activate overlapping parameter subsets.

Backward Pass: Direct Feedback Alignment

Rather than transporting gradients through exact transposed forward weights, each layer receives error information through fixed random projections.

Output error
↓
Fixed random feedback
↓
All hidden layers

Learning adapts the forward weights to exploit these fixed feedback pathways.

Function Sculpting

The only adaptive component is the forward computation within atlas-selected charts.

Gradient descent performs **Function Sculpting**:

Parameter Atlas
selects where computation occurs

Function Sculpting
learns what computation occurs

For a fixed input, the network implements:

$$y = M(x)x$$

where the effective mapping $M(x)$ is assembled from the atlas-selected parameter subsets.

Learning therefore sculpts a family of local mappings over the charts of the atlas.

Fixed Computational Geography Hypothesis

The architecture can be summarized as:

Input
↓
Fixed computational geography
(Parameter Atlas)

↓
Adaptive local computation
(Function Sculpting)

↓
Output

Error
↓
Fixed credit geography
(DFA)

The central hypothesis is:

> Exact gradient transport and learned recursive routing may not be fundamental requirements for effective learning.

Instead, a stable computational and credit-assignment geography may suffice.

Potential Advantages

Reduced Co-Adaptation

Neither routing nor feedback pathways evolve during training.

This removes several sources of instability.

Improved Parameter Utilization

Atlas routing may distribute updates more uniformly across parameters.

Increased Training Stability

Fixed routing and fixed feedback may reduce optimization feedback loops.

Improved Interpretability

The computational geography is explicit and fixed.

Analysis may focus on:

- chart occupancy;
- local function smoothness;
- parameter usage statistics;
- chart adjacency graphs.

Potential Failure Modes

Atlas Slot-Machine Effect

Small input perturbations may alter the selected chart:

small input change
→ different chart
→ different route
→ different effective mapping

Potentially harming generalization.

Routing Chaos in Conventional Networks

Recursive routing architectures may instead exhibit:

small input change
→ altered early routing
→ altered later routing
→ compounding pathway divergence

The relative severity of these two failure modes remains unknown.

Experimental Questions

1. Can a fixed Parameter Atlas combined with DFA match backpropagation-based networks?
2. Does fixed routing improve parameter utilization?

3. Do conventional networks exhibit measurable routing chaos?
4. Is atlas-induced instability less severe than recursively amplified routing divergence?
5. Can stable computational geographies improve learning robustness?

Broader Interpretation

This proposal suggests a radical simplification:

> Neural computation may consist primarily of local function sculpting performed within a fixed computational geography.

If true, much of the complexity of modern architectures may reflect implementation choices rather than fundamental necessities.

Keywords

locality-sensitive hashing, Direct Feedback Alignment, sparse neural networks, conditional computation, manifold hypothesis, routing chaos, fixed computational geography, random projections, mixture of experts, sparse routing, effective linearization.